

# HiMAP: Learning Heuristics-Informed Policies for Large-Scale Multi-Agent Pathfinding

Extended Abstract

Huijie Tang\*  
KAIST  
Daejeon, South Korea  
raylan.tang@kaist.ac.kr

Federico Berto\*  
KAIST  
Daejeon, South Korea  
fberto@kaist.ac.kr

Zihan Ma  
KAIST  
Daejeon, South Korea  
zihanma@kaist.ac.kr

Chuanbo Hua  
KAIST  
Daejeon, South Korea  
cbhua@kaist.ac.kr

Kyuree Ahn  
Omelet  
Daejeon, South Korea  
ahnkyuree@gmail.com

Jinkyoo Park  
KAIST, Omelet  
Daejeon, South Korea  
jinkyoo.park@kaist.ac.kr

## ABSTRACT

Large-scale multi-agent pathfinding (MAPF) presents significant challenges in several areas. As systems grow in complexity with a multitude of autonomous agents operating simultaneously, efficient and collision-free coordination becomes paramount. Traditional algorithms often fall short in scalability, especially in intricate scenarios. Reinforcement Learning (RL) has shown potential to address the intricacies of MAPF; however, it has also been shown to struggle with scalability, demanding intricate implementation, lengthy training, and often exhibiting unstable convergence, limiting its practical application. In this paper, we introduce Heuristics-Informed Multi-Agent Pathfinding (HiMAP), a novel scalable approach that employs imitation learning with heuristic guidance in a decentralized manner. We train on small-scale instances using a heuristic policy as a teacher that maps each single agent observation information to an action probability distribution. During pathfinding, we adopt several inference techniques to improve performance. With a simple training scheme and implementation, HiMAP demonstrates competitive results in terms of success rate and scalability in the field of imitation-learning-only MAPF, showing the potential of imitation-learning-only MAPF equipped with inference techniques.

## KEYWORDS

Multi-Agent Systems; Pathfinding; Heuristics Guidance; Imitation Learning

### ACM Reference Format:

Huijie Tang\*, Federico Berto\*, Zihan Ma, Chuanbo Hua, Kyuree Ahn, and Jinkyoo Park. 2024. HiMAP: Learning Heuristics-Informed Policies for Large-Scale Multi-Agent Pathfinding: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 3 pages.

\*Equal contributions.



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

## 1 INTRODUCTION

Multi-Agent Pathfinding (MAPF) entails finding collision-free paths for a group of agents while minimizing makespan or flowtime [14]. MAPF is a domain with several practical applications in our increasingly automated world, including warehouse robotics [16], aviation [10], and digital gaming [7] scenarios.

Centralized strategies have offered solutions by reducing MAPF to other familiar problems, e.g., ILP [17] and SAT [15], or employing search-based algorithms, e.g., Conflict-Based Search (CBS) [13] and its improved versions ECBS [3] and EECBS [6]. Yet, scalability remains hard to obtain, especially as the number of agents increases.

Reinforcement Learning (RL) approaches [2, 5, 8, 9, 12] offer another way to solve the MAPF problem by treating it as a sequential decision-making problem rather than a centralized planning problem. However, this approach often involves intricate implementation and lengthy training and often exhibits unstable convergence.

**Contributions.** In this paper, we propose Heuristics-Informed Multi-Agent Pathfinding (HiMAP), a novel method that relies only on imitation learning (IL) through which the neural network model learns to act from heuristics expert solutions. Our contributions include: 1) We formulate the MAPF problem as an imitation learning problem where the model learns to choose actions by imitating expert decisions. 2) We propose a simple yet effective training scheme based on imitation learning of small-scale solutions from heuristics solvers. 3) We introduce several inference techniques to improve the performance. 4) We demonstrate the potential of imitation-learning-only MAPF equipped with inference techniques.

## 2 HiMAP

HiMAP trains with the help of heuristics guidance in IL manner. During pathfinding, HiMAP first gathers observations from all agents, where each observation comprises four distinct channels. Subsequently, these observations are fed into the neural network. Leveraging four distinct inference techniques, the neural network produces action for all agents. As a result, the locations of all agents are updated, facilitating a transition to the subsequent state.

**Network Architecture and Observation Space.** We adopt the same neural network as PRIMAL [12], while eliminating the Value and Blocking output head. Similar to PRIMAL, we use a partially observable grid world as the environment. Each agent  $i$  at location  $v_i^t$  has a four-channel observation  $o_i^t$  in a  $9 \times 9$  Field of View (FOV)

centered around itself, plus a  $3 \times 1$  vector pointing from its own position to its goal. At each time step  $t$ , agents act in the grid world either to move to an adjacent vertex or stay at the current vertex.

**Heuristic Guidance.** We train the model with the help of heuristic expert guidance in an imitation learning (IL) manner. We use EECBS [6] to generate a series of expert collision-free paths on randomly generated environments. In the generated heuristic expert paths, each location of a single agent  $i$  at time  $t$ , i.e.,  $v_i^t$ , can be viewed as a decision point, and the decision  $a_i^{t*}$  according to the expert paths is either up, down, left, right, or stay still. HiMAP is trained by minimizing the MSE loss between the model output  $\hat{a}_i^t$  given the input observation  $o_i^t$ , and the expert action decision  $a_i^{t*}$ .

**Inference Techniques.** 1) *Preventing Re-Visit:* During the pathfinding, we find that sometimes agents cannot plan well due to their limited FOV. This will leave some agents oscillating around dead ends and never reaching their goals. Hence, we introduce a history recording scheme that records the history of visited locations and prevents the agents from repetitively re-visiting visited locations during the last few time steps.


2) *Softmax Temperature Adaptation:* We found that planning with only greedy inference does not perform well. Therefore, we adjust the softmax temperature of the final layer of the model such that it can increase or decrease the amount of exploration. During training, the temperature is set to  $\tau = 1$ . During pathfinding, we can adjust the value of  $\tau$  to induce more or less exploration of the map, which can be beneficial for the agent to get unstuck.

3) *Treating Completed Agents as Obstacles:* In this work, we adopt the *stay at target* setting, as defined in [14]. Thus, agent overlapping with its goal can be effectively treated as an obstacle from the model’s perspective. Therefore, we propose changing the map representation for completed agents into obstacles to improve the model performance effectively. We observe that results are improved compared with the situation without this modification, especially with many agents.

4) *Conflict-Free Planning:* We consider four types of conflicts: *swapping conflict* [14], *vertex conflict* [14], conflict with static obstacles, and out-of-bound conflict. During action sampling, the sampled action may guide agents to take invalid actions, leading to one of the conflicts. We implement a two-stage technique to correct invalid actions to enable conflict-free planning. Before sampling, we identify the actions that may result in agents colliding with obstacles or going out of bounds and mask these actions by setting their probability to 0 so they do not influence decision-making. After sampling, the sampled actions may still result in *swapping conflict* or *vertex conflict*. Similar to [8], we employ simple rules to avoid these two conflicts. If these two conflicts happen, the related agents’ states are recursively recovered until no conflicts exist.

### 3 EXPERIMENTS

**Training Data Generation.** We generate 100 random square maps with sizes  $40 \times 40$  and obstacle density 0.3, and another 100 maps with sizes  $80 \times 80$  and the same obstacle density. For each map, we create five *scenario files* [14], each containing 4, 8, 16, 32, or 64 start-goal location pairs, resulting in a total of 1000 *scenario files*. We solve all 1000 Multi-Agent Path Finding (MAPF) problems using EECBS with a suboptimality factor of 1.2. The training dataset



**Figure 1: Success rate of HiMAP and ablated models. Baseline: HiMAP without preventing re-visit (No History) and treating completed agents as obstacles (No TCAO).**

consists of pairs  $(o_i^t, a_i^{t*})$ , generated from expert solutions. Generating all expert paths takes less than a minute on an INTEL CORE i7-6850K. Other hyperparameters include a learning rate starting at  $5 \times 10^{-5}$ , decreasing by 80% every eight epochs, and a total of 35 training epochs. Training takes approximately 100 minutes on a single NVIDIA TITAN X GPU with 12GB of VRAM. Reproducible source code is available on GitHub<sup>1</sup>.

We follow [4] and define the success rate as the average percentage of agents reaching the destination without collisions out of the total number of agents within the maximum allowed time steps. The maximum allowed time step during testing is 256 for the  $40 \times 40$  map and 386 for the  $80 \times 80$  map. For pathfinding, we use history size  $H = 5$ , i.e., the agent is banned from re-visiting the location it has visited in the last five time steps, and softmax temperature  $\tau = 2$ .

The success rate is shown in Fig. 1, demonstrating the power of HiMAP with pure IL equipped with inference techniques. We also show the success rate of ablated models to verify the effectiveness of proposed inference techniques. We believe that the results can be further improved by using a larger training dataset, which will be left for future work.

### 4 DISCUSSION AND CONCLUSION

To the best of our knowledge, [1] is the only paper that employs pure IL to train neural network solvers for MAPF, using a dataset based on single-agent shortest paths [11]. However, these paths do not guarantee collision-free paths for agents or allow agents to observe each other’s behavior. In contrast, equipped with inference techniques, HiMAP uses collision-free paths generated by EECBS, which better represents the MAPF solution in dynamic settings.

We propose HiMAP, an IL approach that employs heuristics expert paths to solve MAPF problems, whose implementation and training are easier and less costly compared with RL-based solvers. The training dataset only contains paths on relatively small-scale MAPF problems, making it easy to collect expert paths with existing heuristics solvers. We further introduced several inference techniques to boost performance during evaluation. We demonstrate HiMAP’s success rate in different environments with high obstacle density, showcasing the potential of imitation-learning-only MAPF equipped with inference techniques.

<sup>1</sup><https://github.com/kaist-silab/HiMAP>

## REFERENCES

- [1] Natalie Abreu. 2022. Efficient Deep Learning for Multi Agent Pathfinding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 13122–13123.
- [2] Kyuree Ahn, Heemang Park, and Jinkyoo Park. 2023. Learning to Schedule in Multi-Agent Pathfinding. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 7326–7332.
- [3] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. 2014. Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Proceedings of the International Symposium on Combinatorial Search*, Vol. 5. 19–27.
- [4] Jaehoon Chung, Jamil Fayyad, Younes Al Younes, and Homayoun Najjaran. 2023. Learning to Team-Based Navigation: A Review of Deep Reinforcement Learning Techniques for Multi-Agent Pathfinding. *arXiv preprint arXiv:2308.05893* (2023).
- [5] Mehul Damani, Zhiyao Luo, Emerson Wenzel, and Guillaume Sartoretti. 2021. PRIMAL<sub>2</sub>: Pathfinding via reinforcement and imitation multi-agent learning-lifelong. *IEEE Robotics and Automation Letters* 6, 2 (2021), 2666–2673.
- [6] Jiaoyang Li, Wheeler Ruml, and Sven Koenig. 2021. Eecbs: A bounded-suboptimal search for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 12353–12362.
- [7] Hang Ma, Jingxing Yang, Liron Cohen, TK Kumar, and Sven Koenig. 2017. Feasibility study: Moving non-homogeneous teams in congested video game environments. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 13. 270–272.
- [8] Ziyuan Ma, Yudong Luo, and Hang Ma. 2021. Distributed heuristic multi-agent path finding with communication. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 8699–8705.
- [9] Ziyuan Ma, Yudong Luo, and Jia Pan. 2021. Learning selective communication for multi-agent path finding. *IEEE Robotics and Automation Letters* 7, 2 (2021), 1455–1462.
- [10] Robert Morris, Corina S Pasareanu, Kasper Søe Luckow, Waqar Malik, Hang Ma, TK Satish Kumar, and Sven Koenig. 2016. Planning, Scheduling and Monitoring for Airport Surface Operations.. In *AAAI Workshop: Planning for Hybrid Systems*. 608–614.
- [11] Jingyao Ren, Vikraman Sathiyarayanan, Eric Ewing, Baskin Senbaslar, and Nora Ayanian. 2021. MAPFAST: A deep algorithm selector for multi agent path finding using shortest path embeddings. *arXiv preprint arXiv:2102.12461* (2021).
- [12] Guillaume Sartoretti, Justin Kerr, Yunfei Shi, Glenn Wagner, TK Satish Kumar, Sven Koenig, and Howie Choset. 2019. Primal: Pathfinding via reinforcement and imitation multi-agent learning. *IEEE Robotics and Automation Letters* 4, 3 (2019), 2378–2385.
- [13] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence* 219 (2015), 40–66.
- [14] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. 2019. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the International Symposium on Combinatorial Search*, Vol. 10. 151–158.
- [15] Pavel Surynek, Ariel Felner, Roni Stern, and Eli Boyarski. 2016. Efficient SAT approach to multi-agent path finding under the sum of costs objective. In *Proceedings of the twenty-second european conference on artificial intelligence*. 810–818.
- [16] Peter R Wurman, Raffaello D’Andrea, and Mick Mountz. 2008. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine* 29, 1 (2008), 9–9.
- [17] Jingjin Yu and Steven M LaValle. 2013. Planning optimal paths for multiple robots on graphs. In *2013 IEEE International Conference on Robotics and Automation*. IEEE, 3612–3617.